

Social Trust Prediction via Max-norm Constrained 1-bit Matrix Completion

Jing Wang¹, Jie Shen², Huan Xu³

¹ Hefei University of Technology, China
wangjing@mail.hfut.edu.cn

² Rutgers University, USA
js2007@rutgers.edu

³ National University of Singapore, Singapore
mpexuh@nus.edu.sg

Abstract

Social trust prediction addresses the significant problem of exploring interactions among users in social networks. Naturally, this problem can be formulated in the matrix completion framework, with each entry indicating the trustness or distrustness. However, there are two challenges for the social trust problem: 1) the observed data are with sign (1-bit) measurements; 2) they are typically sampled non-uniformly. Most of the previous matrix completion methods do not well handle the two issues. Motivated by the recent progress of max-norm, we propose to solve the problem with a 1-bit max-norm constrained formulation. Since max-norm is not easy to optimize, we utilize a reformulation of max-norm which facilitates an efficient projected gradient decent algorithm. We demonstrate the superiority of our formulation on two benchmark datasets.

1 Introduction

With the increasing popularity of social networks, there exist many interesting and difficult problems, such as friends recommendation, information propagation, etc. In this paper, we study the problem of social trust prediction, which aims to estimate the positive or negative relationship among the users based on the existing trustness information associated with them. This problem plays an important role in social networks as the system can block the invitation from someone that the user does not trust, or recommend new friends who enjoys a high reputation.

Naturally, the social trust problem can be formulated within the matrix completion framework [Liben-Nowell and Kleinberg, 2007] [Billsus and Pazzani, 1998]. That is, the (i, j) -th entry of the observed data matrix $Z \in \mathbb{R}^{n \times n}$ is a 1-bit code implying that the i -th user trusts the j -th user if $Z_{ij} = 1$. Here, n denotes the number of users. However, what we observe is only a small fraction of the entries, whose values are zero. And our goal is to estimate the missing entries according to the 1-bit measurements in Z .

Note that the problem is ill-posed if no assumption is imposed on the structure of the data. To solve the problem,

a number of methods are proposed. Generally, existing social trust prediction methods fall into three categories. The first category is based on similarity measures or the structural context similarity [Newman, 2001] [Chowdhury, 2010] [Katz, 1953] [Jeh and Widom, 2002], motivated by the intuition that an individual tends to trust their neighbors, or the ones with similar trusted people. The second is based on low rank matrix completion [Billsus and Pazzani, 1998] [Cai *et al.*, 2010] [Huang *et al.*, 2013], which assumes that the underlying data matrix is low-rank or can be approximate by a low-rank matrix. The third one models the problem as a binary classification one and utilizes techniques such as logistic regression [Leskovec *et al.*, 2010].

Challenges. However, there are two issues emerging in social trust prediction which are not well characterized by the algorithms in previous works. First, the value of the observed entry is either 1 or -1 , which is analogous to the binary classification problem. But in our problem, we are handling much more complex matrix data. Fortunately, [Srebro *et al.*, 2004] presented a maximum margin matrix factorization framework that unifies the binary problem for vector case and matrix case. The key idea in their work is a low-norm matrix factorization, which will also be utilized in this paper. Second, the locations of the entries are sampled non-uniformly, which gaps the theory and practice for a lot of matrix completion algorithms. To tackle this challenge, we suggest using the max-norm as a convex surrogate for the rank function, which is shown to be superior to the well-known nuclear norm when addressing the non-uniform data [Salakhutdinov and Srebro, 2010].

Our contributions are two-folds: 1) To the best of our knowledge, we are the first to address the social trust prediction problem by utilizing a max-norm constrained formulation. 2) Although a max-norm constrained problem can be solved by SDP solvers and an accurate enough solution can be achieved, we here utilize a projected gradient algorithm that is scalable to large scale datasets. We empirically show the improvement of our formulation for the non-uniform 1-bit benchmarks compared to state-of-the-art solvers.

2 Related Work

Social interaction is investigated intensively in the last decades. The social interaction indicates the friendship, support, enemy or disapproval as shown in Figure 1. Online users




Figure 1: Illustration of the data matrix of social trust. Each column of the data matrix is a rating sample. Each entry is a tag that a user assigns to another user. The symbol “+” denotes the “trust” relationship, “-” denotes “distrust” and “?” means unknown relationship (that we aim to predict). Most of the relationships are unknown, so the data is sparse. Typically, each individual has own preference and friendship network, making the data non-uniform. Also noted that the data matrix may not be symmetric.

rely on the trustworthiness information to filter information, establish collaboration or build social bonds. Social networks rely on the trust information to make recommendation, attract users from other circles, or influencing public opinions. Thus, the exploration of social trust has a wide range of applications, and has emerged as an important topic in social network research. A number of methods are proposed.

One kind of methods are based on the similarity measurement. Specifically, Jaccard’s coefficient is commonly used to measure the probability two items that have a relationship. Inspired by the metric, Jeh and Widom [Jeh and Widom, 2002] proposed a domain-independent similarity measurement, SimRank. [Newman, 2001] directly defined a score to verify the correlation between common neighbors. Some methods are based on relational data modeling, structural proximity measures and stochastic relational model [Getoor and Diehl, 2005] [Liben-Nowell and Kleinberg, 2007] [Yu *et al.*, 2009]. The above mentioned methods are mainly derived from the solutions of link prediction. The link prediction is oriented to network-level prediction, whereas social trust prediction focuses on person-level. Another class of methods are derived from the collaborative filtering methods, such as clustering techniques [Sarwar *et al.*, 2001], model-based methods [Hofmann and Puzicha, 1999], and the matrix factorization models [Srebro *et al.*, 2003] [Mnih and Salakhutdinov, 2007]. However, the data matrix of trust has some structure properties different from the

user-item matrix, such as transitivity. Meanwhile, the social trust in reality is extremely sparse. For instance, Facebook has hundreds of millions of users, but most of them have less than 1,000 friends. Besides, the people with similar personality tend to behave similarly. To sum up, the data matrix of social trust has both sparse and low-rank structure. Thus, the social trust prediction problem is especially suitable for the matrix completion model. That is the focus of our paper.

The problem of matrix completion is to recover a low-rank matrix from a subset of entries [Candès and Recht, 2009], which is given by:

$$\begin{aligned} \min_X \quad & \text{rank}(X) \\ \text{s.t.} \quad & P_\Omega(X) = P_\Omega(Z) \end{aligned} \quad (2.1)$$

where $Z \in \mathbb{R}^{p \times n}$ is the data matrix, X is the recovered matrix, and Ω is the index set of observed entries. The optimization problem (2.1) is not only NP-hard, but requires double exponential time complexity with the number of samples [Recht *et al.*, 2010]. To solve the above problem, one alternative is to use nuclear norm as a relaxation to the rank function:

$$\begin{aligned} \min_X \quad & \|X\|_* \\ \text{s.t.} \quad & P_\Omega(X) = P_\Omega(M) \end{aligned} \quad (2.2)$$

where $\|X\|_*$ denotes the sum of singular values of matrix X . [Cai *et al.*, 2010] developed a first-order procedure to solve the convex problem (2.2), namely singular value thresholding (SVT). [Jain *et al.*, 2010] minimized the rank minimization by the singular value projection (SVP) algorithm. [Keshavan *et al.*, 2010] solved the problem by first trimming each row and column with too few entries, then compute the truncated SVD of the trimmed matrix. Under certain conditions, it showed accurate recovery on the order of $nd \log n$ samples (n is the number of samples, d is the rank of recovered matrix). With the rapid development of matrix completion problem, some more efficient methods have been proposed [Candès and Tao, 2010][Gross, 2011][Wang and Xu, 2012][Huang *et al.*, 2012].

However, all the methods mentioned above use the nuclear norm as the surrogate to the rank, whose exact recovery can be guaranteed only when the data are sampled uniformly, which is not practical in real world applications. On the other hand, recent empirical on max-norm [Srebro *et al.*, 2004] shows promising results for non-uniform data if one utilize the max-norm as a surrogate [Salakhutdinov and Srebro, 2010]. Notably, for some specific problems, such as collaborative filtering, [Srebro and Shraibman, 2005] proved that the generalization error bound for max-norm is better than the nuclear norm. More recently, [Shen *et al.*, 2014] reported encouraging results on the subspace recovery task (which is closely relevant to matrix completion). Since the social trust data is non-uniformly sampled, we believe that a max-norm regularized formulation can better handle the challenge than the nuclear norm. Our formulation is also inspired by a recent theoretical study on matrix completion with 1-bit measurement [Cai and Zhou, 2013], which established a minimax

lower bound on the general sampling model and derived the optimal convergence rate in terms of Frobenius norm loss. Furthermore, there are several practical algorithms to solve max-norm regularized or max-norm constrained problems, see [Lee *et al.*, 2010] and [Shen *et al.*, 2014] for example.

2.1 Overview

After review of related work in Section 2, we introduce the notations and formulate the problem in Section 3. Then we give algorithm to solve the max-norm constrained 1-bit matrix completion (MMC) problem in Section 4. Meanwhile, we also provide an equivalent SDP formulation for the MMC, which can be accurately solved at the expense of efficiency. Then we report the empirical study on two benchmark datasets in Section 5. Section 6 concludes this paper and discusses possible future work.

3 Notations and Problem Setup

In this section, we introduce the notations that will be used in this paper. Capital letters such as M are used for matrices and lowercase bold letters such as \mathbf{v} denotes vectors. The i -th row and j -th column of a matrix M is denoted by $\mathbf{m}(i)$ and \mathbf{m}_j respectively, and the (i, j) -th entry is denoted by M_{ij} . For a vector \mathbf{v} , we use v_i to denote its i -th element. We denote the ℓ_2 norm of a vector \mathbf{v} by $\|\mathbf{v}\|_2$. For a matrix $M \in \mathbb{R}^{p \times n}$, we denote the Frobenius norm by $\|M\|_F$ and $\|M\|_{2,\infty}$ denotes the maximum ℓ_2 row norm of M , i.e.,

$$\|M\|_{2,\infty} := \max_{i=1}^p \|\mathbf{m}(i)\|_2.$$

We further define the *max-norm* of M [Linial *et al.*, 2007],

$$\|M\|_{\max} = \min_{U, V, M=UV^\top} \max\{\|U\|_{2,\infty}^2, \|V\|_{2,\infty}^2\}, \quad (3.1)$$

where we enumerate all possible factorizations to obtain the minimum.

Intuition on max-norm. At a first sight, the max-norm is hard to understand. We simply explain why it is a tighter approximation to the rank function than the nuclear norm. Again, we write the nuclear norm of M as a factorization form [Recht *et al.*, 2010]:

$$\|M\|_* := \min_{U, V, M=UV^\top} \frac{1}{2} (\|U\|_F^2 + \|V\|_F^2).$$

Note that the Frobenius norm is the sum of the square of the ℓ_2 row norm. Thus, a nuclear norm regularizer actually constrains the average of the ℓ_2 row norm, while the max-norm constrains the maximum of the ℓ_2 row norm!

Given the observed data $Z \in \mathbb{R}^{p \times n}$, we are interested in approximating Z with a low-rank matrix X , which can be formulated by,

$$\begin{aligned} & \text{minimize } \frac{1}{2} \|\mathcal{P}_\Omega(Z - X)\|_F, \\ & \text{s.t. } \text{rank}(X) \leq d, \end{aligned}$$

where Ω is an index set of observed entries and d is some expected rank. $\mathcal{P}_\Omega(M)$ is a projection operator on a matrix M such that $\mathcal{P}_\Omega(m_{ij}) = m_{ij}$ if $(i, j) \in \Omega$ and zero

otherwise. However, it is usually intractable to optimize the above program as the rank function is non-convex and non-continuous [Recht *et al.*, 2010]. One common approach is to use the nuclear norm as a convex surrogate to the rank function. However, it is well known that the nuclear norm cannot well handle the non-uniform data. Motivated by the recent progress in max-norm [Salakhutdinov and Srebro, 2010; Cai and Zhou, 2013; Shen *et al.*, 2014], we use the max-norm as an alternative convex relaxation, which gives the following formulation:

$$\begin{aligned} & \min_X \frac{1}{2} \|\mathcal{P}_\Omega(Z - X)\|_F^2, \\ & \text{s.t. } \|X\|_{\max} \leq \lambda^2, \end{aligned} \quad (3.2)$$

where λ is some tunable parameter.

4 Algorithm

The max-norm is convex and moreover, it can be solved by any SDP solver. Formally, we have the following lemma:

Lemma 4.1 ([Srebro *et al.*, 2004]). *For any matrix $X \in \mathbb{R}^{p \times n}$ and $\lambda \in \mathbb{R}$, $\|X\|_{\max} \leq \lambda$ if and only if there exist $A \in \mathbb{R}^{p \times p}$ and $B \in \mathbb{R}^{n \times n}$, such that $\begin{bmatrix} A & X \\ X^\top & B \end{bmatrix}$ is semi-definite positive and each diagonal element in A and B is upper bounded by λ .*

With Lemma 4.1 on hand, one can formulate Problem 3.2 as an SDP:

$$\begin{aligned} & \min_{X, A, B} \frac{1}{2} \|\mathcal{P}_\Omega(Z - X)\|_F^2, \\ & \text{s.t. } A_{ii} \leq \lambda^2, B_{jj} \leq \lambda^2, \forall i \in [p], j \in [n], \\ & \quad \begin{bmatrix} A & X \\ X^\top & B \end{bmatrix} \succeq 0. \end{aligned} \quad (4.1)$$

And this program can be solved by any SDP solver to obtain accurate enough solution.

However, SDP solvers are not scalable to large matrices. Thus, in this paper, we apply a projected gradient method to solve Problem (3.2), which is due to [Lee *et al.*, 2010]. A key technique is the reformulation of the max-norm (3.1). Assume that the rank of the optimal solution X^* produced by the SDP (4.1) is at most d . Then we can safely factorize $X = UV^\top$, with $U \in \mathbb{R}^{p \times d}$ and $V \in \mathbb{R}^{n \times d}$. Combining the factorization and the definition, we obtain the following equivalent program :

$$\begin{aligned} & \min_{U, V} \frac{1}{2} \|\mathcal{P}_\Omega(Z - UV^\top)\|_F^2, \\ & \text{s.t. } \|U\|_{2,\infty} \leq \lambda, \|V\|_{2,\infty} \leq \lambda. \end{aligned} \quad (4.2)$$

Note that the gradient of the objective function w.r.t. U and V can be easily computed. That is,

$$\begin{aligned} \nabla_U f(Z, U, V) &= \mathcal{P}_\Omega((UV^\top - Z)V), \\ \nabla_V f(Z, U, V) &= \mathcal{P}_\Omega((VU^\top - Z^\top)U). \end{aligned} \quad (4.3)$$

Here, for simplicity we define

$$f(Z, U, V) = \frac{1}{2} \|\mathcal{P}_\Omega(Z - UV^\top)\|_F^2.$$

Algorithm 1 Max-norm Constrained 1-Bit Matrix Completion (MMC)

Input: $Z \in \mathbb{R}^{p \times n}$ (observed samples), parameters λ , initial solution (U_0, V_0) , maximum iteration τ .

Output: Optimal solution (U^*, V^*) .

1: **for** $t = 1$ to τ **do**

2: Compute the gradient by Eq. (4.3):

$$U'_t = \nabla_U f(Z, U, V_{t-1})|_{U=U_{t-1}},$$

$$V'_t = \nabla_V f(Z, U_{t-1}, V)|_{V=V_{t-1}}.$$

3: Compute the step size α_t according to Armijo rule.

4: Compute the new iterate:

$$U_t = \Pi(U_{t-1} - \alpha_t U'_t),$$

$$V_t = \Pi(V_{t-1} - \alpha_t V'_t).$$

5: **end for**

The inequality constraints can be addressed by a projection step. That is, when we have a new iterate (U_t, V_t) at the t -th iteration, we can check if they violate the constraints. If not, we can proceed to the next iteration. Otherwise, we can scale the rows of U and/or V by $\frac{\lambda}{\|U\|_{2,\infty}}$ and/or $\frac{\lambda}{\|V\|_{2,\infty}}$ respectively. In this way, we have the projection operator:

$$\Pi(M) = \begin{cases} \frac{\lambda}{\|M\|_{2,\infty}} M, & \text{if } \|M\|_{2,\infty} > \lambda, \\ M, & \text{otherwise.} \end{cases} \quad (4.4)$$

If we further pick the step size α_t via the Armijo rule [Armijo and others, 1966], it can be shown that the sequence of (U_t, V_t) will converge to a stationary point [Bertsekas, 1999]. The algorithm is summarized in Algorithm 1.

The benefits of applying the factorization on the max-norm are two-folds: 1) the memory cost can be significantly reduced from $O(pn)$ of SDP to $O(d(p+n))$. 2) it facilitates the projected gradient algorithm, which is computationally efficient when working on large matrices (see Section 5). However, note that Problem (4.2) is non-convex. Fortunately, [Burer and Monteiro, 2005] proved that as long as we pick a sufficiently large value for d , then any local minimum of Eq. (4.2) is a global optimum. In Section 5, we will report the influence of d on the performance. Actually, in Algorithm 1, the stopping criteria is set to be a maximum iteration. One may also check if it reaches a local minima as the stopping criteria, as discussed in [Cai and Zhou, 2013].

4.1 Heuristic on λ

The λ is the only tunable parameter in our algorithm. For our problem, note that the data is of 1-bit measurements, *i.e.*, $|Z_{ij}| = 1$ for $(i, j) \in \Omega$. Also note that $Z_{ij} = \mathbf{u}(i)\mathbf{v}(j)^\top$. Thus, $|\mathbf{u}(i)\mathbf{v}(j)^\top| = 1$. So we have $\lambda \geq 1$. However, if we choose a large λ , the estimation $|X_{ij}|$ may deviate away from 1. We find that $\lambda = 1.2$ lead to satisfactory improvement.

5 Experiments

In this section, we empirically evaluate our method for the matrix completion performance. We will first introduce the used datasets. In the experimental settings, we present the comparative methods and evaluation metrics. Then we report encouraging results on two benchmark datasets. We also examine the influence of matrix rank d .

5.1 Datasets

We conduct the experiments on two benchmark datasets: Epinions and Slashdot. In these two datasets, the users are connected by explicit positive (trust) or negative (distrust) links (*i.e.*, the 1-bit measurements in Z). The first dataset contains 119,217 nodes (users) and 841,000 edges (links), 85.0% of which are positive. The Slashdot dataset contains 82,144 users and 549,202 links, and 77.4% of the edges are labeled as positive. Table 1 gives a summary description about the subset used in our experiment.

It is clear that the distribution of links are not uniform since each user has his/her individual preference and own friendship network. Following [Huang *et al.*, 2013], we select 2,000 users with the highest degrees from each dataset to form the observation matrix Z .

Table 1: Description of 2 datasets

Dataset	Epinions	Slashdot
#of Users	2,000	2,000
#of Trust	171,731	68,932
#of Distrust	18,916	20,032

5.2 Experimental Settings

Baselines. We choose four state-of-the-art methods as baselines, including SVP [Jain *et al.*, 2010], SVT [Cai *et al.*, 2010], OPTSpace [Keshavan *et al.*, 2010] and RRMC [Huang *et al.*, 2013]. Since SVT and RRMC need a specified rank, we tune the rank for these methods and choose the best performance as the final result.

Evaluation Metric. Let T be the index set of all observed entries. We use two evaluation metrics to measure the performance, mean average error (MAE) and root mean square error (RMSE), computed as follows:

$$MAE = \frac{\sum_{(i,j) \in T \setminus \Omega} (X_{ij} - M_{ij})}{(|T| - |\Omega|)},$$
$$RMSE = \sqrt{\frac{\sum_{(i,j) \in T \setminus \Omega} (X_{ij} - M_{ij})^2}{(|T| - |\Omega|)}},$$

where $|T|$ denotes the cardinality of T .

Training and Testing. We randomly split the dataset for training and testing. In particular, the number of observation measurements Ω for training ranges from 10% to 60%, with step size 10%. For each split, we run all the algorithms for 20 trials, with the training data in each trail being randomly sampled. Then, we report the mean and standard deviation of MAE and RMSE over all 20 trials.

Table 2: MAE Results on Epinions Dataset

Observed entries (%)	Methods				
	SVT	OPTSpace	SVP	RRMC	MMC
10	0.359 \pm 0.004	0.289 \pm 0.019	0.450 \pm 0.008	0.576 \pm 0.001	0.254\pm0.003
20	0.394 \pm 0.022	0.236 \pm 0.005	0.294 \pm 0.002	0.518 \pm 0.002	0.212\pm0.003
30	0.360 \pm 0.057	0.219 \pm 0.009	0.248 \pm 0.001	0.460 \pm 0.002	0.201\pm0.002
40	0.410 \pm 0.099	0.205 \pm 0.008	0.224 \pm 0.001	0.418 \pm 0.002	0.193\pm0.001
50	0.471 \pm 0.129	0.197 \pm 0.007	0.210 \pm 0.001	0.386 \pm 0.001	0.190\pm0.002
60	0.476 \pm 0.146	0.197\pm0.003	0.199 \pm 0.001	0.362 \pm 0.002	0.206 \pm 0.003

Table 3: RMSE Results on Epinions Dataset

Observed entries (%)	Methods				
	SVT	OPTSpace	SVP	RRMC	MMC
10	0.513 \pm 0.010	0.530 \pm 0.021	0.610 \pm 0.010	0.650 \pm 0.001	0.466\pm0.004
20	0.559 \pm 0.031	0.456 \pm 0.005	0.459 \pm 0.002	0.606 \pm 0.002	0.406\pm0.004
30	0.532 \pm 0.082	0.422 \pm 0.011	0.415 \pm 0.002	0.563 \pm 0.002	0.383\pm0.002
40	0.620 \pm 0.171	0.406 \pm 0.015	0.394 \pm 0.002	0.533 \pm 0.002	0.371\pm0.002
50	0.719 \pm 0.225	0.398 \pm 0.016	0.381 \pm 0.001	0.509 \pm 0.001	0.364\pm0.001
60	0.728 \pm 0.288	0.403 \pm 0.009	0.371 \pm 0.002	0.491 \pm 0.002	0.365\pm0.002

5.3 Experimental Results

We report detailed results from Table 2 to Table 5. From the results in Tables 2 and 3, we observe that MMC outperforms the other methods in terms of both evaluation metrics on the Epinions dataset most of the time. In particular, when there are few observations 10 % (which indicates a hard task), MMC obtains the RMSE of 0.466, much better than OPTSpace (0.530), SVP (0.610) and RRMC (0.650). Except on the case with 60% observed entries, OPTSpace obtains the smallest MSE with 0.197, but our algorithm is comparative with 0.206. In a nut of shell, the gap between MMC and the baselines becomes larger as the fraction of observed entries decreases.

Similarly, our method achieves the least MAE and RMSE on the Slashdot dataset (see Table 4 and 5). For instance with 30 % observed entries, MMC obtains the MAE with less than 0.4, much better than the comparative methods, such as SVT (0.513), OPTSpace (0.427), SVP (0.501) and RRMC (0.686). In terms of RMSE, in the case of 20% observed entries, the RMSE values of other methods are all above 0.7 while our method reaches 0.679. In sum, our method is superior than the comparative methods on two real-life datasets in terms of MSE and RMSE most of the time.

Since we have studied the effectiveness of our method, here we examine the computational efficiency in Table 6, which is important for practical applications. To test the time complexity of the methods, we report the averaged time cost on the Epinions dataset with 10% observed entries and Slashdot with 20% observed entries. To illustrate the trade-off between accuracy and efficiency, we also report the MAE and RMSE. As we see, SVP is the most efficient method, whose running time is 0.84 seconds on Epinions while ours is 1.92 seconds. On Slashdot, it also achieves the best performance in terms of efficiency. However, our method enjoys a significant improvement of MAE and RMSE compared to all base-


Figure 2: MAE and RMSE in terms of different rank d on the Epinions dataset.

Table 4: MAE Results on Slashdot Dataset

Observed entries (%)	Methods				
	SVT	OPTSpace	SVP	RRMC	MMC
10	0.679±0.008	0.554±0.017	0.755±0.005	0.715±0.001	0.546±0.008
20	0.562±0.004	0.458±0.008	0.582±0.007	0.704±0.001	0.437±0.008
30	0.513±0.030	0.427±0.009	0.501±0.003	0.686±0.003	0.395±0.006
40	0.506±0.041	0.395±0.009	0.460±0.002	0.647±0.006	0.380±0.009
50	0.495±0.060	0.376±0.004	0.432±0.002	0.609±0.002	0.366±0.003
60	0.520±0.065	0.362±0.011	0.413±0.002	0.585±0.002	0.350±0.002

Table 5: RMSE Results on Slashdot Dataset

Observed entries (%)	Methods				
	SVT	OPTSpace	SVP	RRMC	MMC
10	0.788±0.008	0.826±0.021	0.873±0.004	0.829±0.001	0.774±0.008
20	0.718±0.020	0.704±0.013	0.746±0.007	0.821±0.001	0.679±0.009
30	0.680±0.043	0.652±0.006	0.680±0.003	0.807±0.002	0.633±0.008
40	0.670±0.056	0.620±0.009	0.647±0.003	0.778±0.005	0.615±0.011
50	0.642±0.055	0.596±0.009	0.624±0.003	0.749±0.002	0.581±0.006
60	0.699±0.084	0.577±0.011	0.609±0.002	0.730±0.002	0.566±0.002

Table 6: The trade-off between accuracy and efficiency.

Dataset	SVT			OPTSpace			SVP			RRMC			MMC		
	MAE	RMSE	Time	MAE	RMSE	Time	MAE	RMSE	Time	MAE	RMSE	Time	MAE	RMSE	Time
Epinions	0.466	0.618	26.76	0.288	0.531	15.26	0.450	0.610	0.84	0.576	0.651	43.11	0.262	0.481	1.92
Slashdot	0.618	0.728	50.73	0.458	0.705	9.23	0.582	0.746	0.83	0.715	0.830	44.12	0.437	0.679	2.41

lines. Also, our algorithm is orders of magnitude faster than other three baselines. This implies that MMC favors a good trade-off between the accuracy and efficiency.

5.4 Examine The Influence of d

The non-convex reformulation (4.2) requires an explicit rank estimation d on the true matrix. In this section, we investigate the influence of d on the Epinions dataset as an example. The rank d is chosen from [1, 5, 50, 100, 300, 500] and the results are plot in Figure 2. We observe that the rank has little influence on the performance. This is possibly because that the actual data has a low-rank structure (close to rank one). And from [Burer and Monteiro, 2005], we know that if d is larger than the actual rank, any local minimum of Eq. (4.2) is also a global optima.

6 Conclusion and Future Work

In this paper, we formulated the social trust prediction in the matrix completion framework. In particular, due to the special structure of the social trust problem, *i.e.*, the measurements are 1-bit and the observed entries are non-uniformly sampled, we presented a max-norm constrained 1-bit matrix completion (MMC) algorithm. Since SDP solvers are not scalable to large scale matrices, we utilized a non-convex reformulation of the max-norm, which facilitates an efficient projected gradient decent algorithm. We empirically examined our algorithm on two benchmark datasets. Compared to other state-of-the-art matrix completion formulations, MMC

consistently outperformed them, which meets with recently developed theories on max-norm. We also studied the trade-off between the accuracy and efficiency and observed that MMC achieved superior accuracy while keeping comparable computational efficiency.

The max-norm has been studied for several years and in many applications, such as collaborative filtering, clustering, subspace recovery. It is empirically and theoretically shown to be superior than the popular nuclear norm. This work investigates the power of max-norm for social trust problem and demonstrates encouraging results. It is interesting and promising to apply max-norm as a convex surrogate to other practical problems such as face recognition, subspace clustering etc.

References

- [Armijo and others, 1966] Larry Armijo et al. Minimization of functions having lipschitz continuous first partial derivatives. *Pacific Journal of mathematics*, 16(1):1–3, 1966.
- [Bertsekas, 1999] Dimitri P Bertsekas. Nonlinear programming. 1999.
- [Billsus and Pazzani, 1998] Daniel Billsus and Michael J Pazzani. Learning collaborative information filters. In *ICML*, volume 98, pages 46–54, 1998.
- [Burer and Monteiro, 2005] Samuel Burer and Renato DC Monteiro. Local minima and convergence in low-rank semidefinite programming. *Mathematical Programming*, 103(3):427–444, 2005.
- [Cai and Zhou, 2013] Tony Cai and Wen-Xin Zhou. A max-norm constrained minimization approach to 1-bit matrix completion. *The Journal of Machine Learning Research*, 14(1):3619–3647, 2013.
- [Cai et al., 2010] Jian-Feng Cai, Emmanuel J Candès, and Zuowei Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.
- [Candès and Recht, 2009] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- [Candès and Tao, 2010] Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- [Chowdhury, 2010] Gobinda Chowdhury. *Introduction to modern information retrieval*. 2010.
- [Getoor and Diehl, 2005] Lise Getoor and Christopher P Diehl. Link mining: a survey. *ACM SIGKDD Explorations Newsletter*, 7(2):3–12, 2005.
- [Gross, 2011] David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.
- [Hofmann and Puzicha, 1999] Thomas Hofmann and Jan Puzicha. Latent class models for collaborative filtering. In *IJCAI*, volume 99, pages 688–693, 1999.
- [Huang et al., 2013] Jin Huang, Feiping Nie, Heng Huang, Yi-Cheng Tu, and Yu Lei. Social trust prediction using heterogeneous networks. *ACM Transactions on Knowledge Discovery from Data*, 7(4):17, 2013.
- [Jain et al., 2010] Prateek Jain, Raghu Meka, and Inderjit S Dhillon. Guaranteed rank minimization via singular value projection. In *NIPS*, pages 937–945, 2010.
- [Jeh and Widom, 2002] Glen Jeh and Jennifer Widom. Simrank: a measure of structural-context similarity. In *ACM KDD*, pages 538–543, 2002.
- [Katz, 1953] Leo Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, 1953.
- [Keshavan et al., 2010] Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *Information Theory, IEEE Transactions on*, 56(6):2980–2998, 2010.
- [Lee et al., 2010] Jason D Lee, Ben Recht, Nathan Srebro, Joel Tropp, and Ruslan R Salakhutdinov. Practical large-scale optimization for max-norm regularization. In *NIPS*, pages 1297–1305, 2010.
- [Leskovec et al., 2010] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. Predicting positive and negative links in online social networks. In *WWW*, pages 641–650, 2010.
- [Liben-Nowell and Kleinberg, 2007] David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *Journal of the American society for information science and technology*, 58(7):1019–1031, 2007.
- [Linial et al., 2007] Nati Linial, Shahar Mendelson, Gideon Schechtman, and Adi Shraibman. Complexity measures of sign matrices. *Combinatorica*, 27(4):439–463, 2007.
- [Mnih and Salakhutdinov, 2007] Andriy Mnih and Ruslan Salakhutdinov. Probabilistic matrix factorization. In *NIPS*, pages 1257–1264, 2007.
- [Newman, 2001] Mark EJ Newman. Clustering and preferential attachment in growing networks. *Physical Review E*, 64(2):025102, 2001.
- [Recht et al., 2010] Benjamin Recht, Maryam Fazel, and Pablo A Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.
- [Salakhutdinov and Srebro, 2010] Ruslan Salakhutdinov and Nathan Srebro. Collaborative filtering in a non-uniform world: Learning with the weighted trace norm. *tc (X)*, 10:2, 2010.
- [Sarwar et al., 2001] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *WWW*, pages 285–295, 2001.
- [Shen et al., 2014] Jie Shen, Huan Xu, and Ping Li. Online optimization for max-norm regularization. In *NIPS*, pages 1718–1726, 2014.
- [Srebro and Shraibman, 2005] Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *Learning Theory*, pages 545–560. 2005.
- [Srebro et al., 2003] Nathan Srebro, Tommi Jaakkola, et al. Weighted low-rank approximations. In *ICML*, volume 3, pages 720–727, 2003.
- [Srebro et al., 2004] Nathan Srebro, Jason DM Rennie, and Tommi Jaakkola. Maximum-margin matrix factorization. In *NIPS*, volume 17, pages 1329–1336, 2004.
- [Wang and Xu, 2012] Yu-Xiang Wang and Huan Xu. Stability of matrix factorization for collaborative filtering. *ICML*, 2012.
- [Yu et al., 2009] Kai Yu, John Lafferty, Shenghuo Zhu, and Yihong Gong. Large-scale collaborative prediction using a nonparametric random effects model. In *ICML*, pages 1185–1192, 2009.